


The Hashed Oct-Tree N-Body Algorithm at a Petaflop

Michael S. Warren, T-2;
Ben Bergen, CCS-7

Cosmological simulations are the cornerstone of theoretical analysis of large-scale structure. During the next few years, projects such as the South Pole Telescope (SPT) and the Dark Energy Survey (DES) will measure the spatial distribution of large-scale structure in enormous volumes of space across billions of years of cosmic evolution. Within a decade, the next generation of observing projects such as the Large Synoptic Survey Telescope (LSST)—in optical wavelengths—and the Square Kilometer Array (SKA)—in the radio spectrum—will gather hundreds to thousands of petabytes of observational data. Advances in modeling must keep pace with observational advances if we are to understand the universe that led to these observations.

We recently demonstrated our hashed oct-tree N-body code (HOT) scaling to 256 k processors on Jaguar at ORNL with a performance of 1.79 Petaflops (single precision) on 2 trillion particles. We also performed preliminary studies with NVIDIA Fermi graphical processing units (GPU), achieving single-GPU performance on our hexadecapole inner loop of 1 Tflop (single precision) and application performance speedup of 2× by offloading the most computationally intensive part of the code to the GPU.

Fig. 1. A halo from a high-resolution simulation containing 134 billion particles. The image is 8 Mpc across, and the spatial resolution of 1 kpc is 1/10th the scale of a pixel in this image.



Understanding the nature of dark matter and dark energy is undoubtedly among the most important unsolved problems in physics. The intrinsic non-linearity of the gravitational evolution of matter in the universe has limited the analytic studies of the problem to small perturbations or restricted symmetries. The only known way to obtain accurate 3D solutions is via numerical simulation. Modern N-body simulation codes provide the theoretical basis for our present understanding of the mass distribution in the universe, and are an essential link in the chain that connects particle physics to cosmology.

Such simulations have been at the forefront of parallel computing since the early 1990s. N-body simulations have grown from 300 particles in 1970 to hundreds of billions of particles today.

We have a long and distinguished history in the development of parallel numerical techniques for solving astrophysical and cosmological N-body problems. We have achieved superior performance on multiple generations of the fastest supercomputers in the world with our HOT, spanning two decades and garnering multiple Gordon Bell Prizes for significant achievement in parallel processing. With a recent Discretionary award on Jaguar at ORNL we have demonstrated our code scaling to 256 k processors with a performance of 1.79 Petaflops (single precision) on 2 trillion particles.

The revolutionary transformation of cosmology from a qualitative to a quantitative science has occurred over just the last 20 years. Driven by a powerful and diverse suite of observations, the parameters describing the large-scale universe are now known to approach 1% precision. This remarkable narrowing of parameters greatly facilitates research on the next prime target of precision cosmology: understanding the growth of cosmological structure in the non-linear regime. Our scientific aims are to use computer simulations to better understand the fundamental properties of the large-scale universe. These questions at the frontier of science include: How do cosmic structures form and evolve? What is dark matter? Why is the universe accelerating?

A great deal of progress has been made over the past 25 years in parallel computing, but much remains the same. One of our first scientific N-body simulations of dark matter in 1990 used 1.1 million particles and was performed on the Caltech/JPL Mark III hypercube in 1990. The 64-node machine used Motorola 68000 microprocessors accelerated with a Weitek floating point co-processor. The simulation was completed in 60 hours, sustaining 160 Mflops with a parallel efficiency of 85%. We have recently demonstrated our ability to use a similar algorithm to perform a simulation on 260 thousand processors with over one trillion (10^{12}) particles, sustaining in excess of 1.5 Petaflops with a parallel efficiency of 90% (Fig. 1). Since our first parallel simulations, the message-passing programming model, time to solution, and parallel efficiency are nearly the same, but the problem size has increased by




Fig. 2. The halo mass function across five orders of magnitude from three 69 billion particle simulations compared with the Tinker (2008) fit.

a factor of a million, and performance by a factor of 10 million. Increasing the performance of an automobile by a factor of 10 million would allow it to travel at the speed of light!

Our simulations are being driven to higher and higher particle numbers due to the simple fact that observations are probing smaller scales to higher accuracy, but we cannot increase our small-scale resolution at fixed particle number without reducing the spatial volume of the simulation,

which then creates errors due to the statistical variance in the large-scale modes. With a recent series of 69-billion-particle simulations performed at LANL, we can quantify these errors at both large and small scales, and calculate the simulation parameters necessary to produce data products that are accurate to 1% or better over the range of interest. An example dark matter halo from one of our high-resolution simulations is shown in Fig. 2.




Fig. 3. The power spectrum of density fluctuations at redshift zero from six of our latest $N=4096^3$ simulations.

by factors of eight from $2.44 \times 10^7 M_\odot$ to $1 \times 10^{11} M_\odot$. The galaxy halo mass function is shown in Fig. 4, comparing it with the fit derived from our earlier work.

Computer simulations enable discovery. In the words of the Astronomy and Astrophysics Decadal Survey: “Through computer modeling, we understand the deep implications of our very detailed observational data and formulate new theories to stimulate further observations.” The only way to accurately model the evolution of dark matter in the universe is through the use of advanced algorithms on massively parallel computers. On the order of a billion dollars will be invested over the coming years in observational projects probing for signatures of dark matter and dark energy. The return on this investment depends a great deal on having a robust and accurate suite of simulations to interpret these observations in the light of our theoretical models. We have demonstrated our ability to achieve the scale and accuracy necessary using our advanced parallel algorithms and a computer with hundreds of Terabytes of memory and Petascale performance.




Fig. 4. Scaling on Jaguar at ORNL measured in June 2012.

Bibliography

Warren, M.S. and Bergen, B., “The Hashed Oct-Tree N -body Algorithm at a Petaflop.” Proceedings IEEE Comp. Soc. (2012).

Tinker, J.L., *Astrophys J* **688**, 709 (2008).